

Remarks

Responsive to applicants' Appeal Brief filed February 6, 2003, prosecution of the application was reopened and the above-referenced Office Action of June 2, 2003 issued.

Applicants have herein elected to respond to this Office Action rather than currently reinstate the appeal in order to more fully develop and address the new issues raised by the Examiner.

In the Office Action, claim 1 is initially rejected under 35 U.S.C. 112, first paragraph, as based on the disclosure which is non-enabling. This rejection is respectfully, but most strenuously, traversed and reconsideration thereof is requested.

To the extent understood, the rejection appears to be based on an allegation that the disclosure does not describe how a node changes topology information. This conclusion is respectfully traversed and is believed counter to the extensive discussion in the specification and drawings of how applicants update and change topology information.

For example, the following discussion is provided on page 14, line 3 - page 18, line 14 of the application as filed:

In accordance with one aspect of the present invention, a reconfiguration protocol is provided for implementing configuration changes to the system. These configuration changes, or reconfigurations, are implemented in each node to reflect changes in the topology of the system. Thus, for example, reconfiguration occurs with the addition or removal of computing resources, such as, for instance, a single node or an entire network of nodes, to the existing system or with, for example, a change in the address or addition or removal of an adapter. Furthermore, since reconfiguration may be necessitated at any time, even in the middle of an executing protocol, the reconfiguration should be made with minimal disturbance to the system and also without interruption to executing protocols.

To implement this particular protocol, information relating to the topology of the configuration, as well as a configuration identifier representing the configuration of the system, is written to the configuration file of FIG. 2 with each event that changes the configuration of the system. As discussed above, this configuration file is then read by each node at the commencement of reconfiguration. After reading the configuration file, each

node stores this information to local memory. The configuration identifier is then utilized to trigger reconfiguration, as discussed in greater detail below.

Referring to FIGS. 9 and 10, reconfiguration is typically initiated by a user, locally, at a particular node, 910 and 1010. As part of a user initiated reconfiguration, new configuration information, including a new configuration identifier is stored to the configuration file. In addition, reconfiguration at a particular node may be triggered by the receipt of a message having a reconfiguration sequence identifier (e.g., a number equal to the maximum possible value of the system) 915 and 1015. In any event, the node that receives the reconfiguration request (from either a user or in the form of a message) responds by reading the configuration file containing the new configuration information, 920 and 1020, and stores this new information to local memory. However, instead of immediately implementing the configuration changes, the node enters a quiescent state for a predetermined period of time, 925 and 1025.

In particular, each node remains in a quiescent state for a period of time sufficient to allow currently running protocols to complete execution under the previous configuration. In addition, the quiescent state is set to allow the propagation of messages containing a reconfiguration sequence identifier to the other nodes of the system, and thus trigger reconfiguration in those other nodes as well, 915 and 1015. Therefore, before the quiescent state terminates at an individual node, all currently running protocols (e.g., death or join protocols) will have completed execution, and all of the other nodes in the system will have received a reconfiguration request and thus will have also entered their own respective quiescent states. Hence, by using a predetermined period of time, each node may terminate its quiescent state without any express communication or acknowledgment from any other nodes.

As discussed above, upon receiving a reconfiguration request, each node enters a quiescent state for a predetermined period of time sufficient to allow currently running protocols to complete execution under the previous configuration. Specifically, each protocol is set so that execution completes after a limited amount of retries are performed with a set period of time between retries. To use the join protocol as an example, a group leader attempting to have a new node commit will transmit only a limited amount of PTC messages, for example, twenty, with, for example, three seconds between the transmission of each PTC message. Thus, in this example, assuming that the join protocol is the lengthiest protocol utilized by the system, the predetermined period of the quiescent state is set to cover, at least, 60 seconds (20 retries X 3 seconds between each retry).

Furthermore, while in this quiescent state, each node also propagates reconfiguration requests to the other nodes of the system (see, 915 and 1030). In accordance with the principles of the invention, the system utilizes a

slightly modified message belonging to another protocol to facilitate the propagation of these reconfiguration requests. More specifically, messages such as, for instance, proclaim, node connectivity, or group connectivity messages are modified to include the aforementioned reconfiguration sequence identifier (e.g., a number equal to the maximum possible value of the system). Additionally, in other embodiments, the frequency of transmission of some messages may be increased to further increase the rate of propagation. Then, upon receipt of these messages by the other nodes of the system, reconfiguration is triggered in those other nodes as well (see, 1015).

To ensure that new configurations are not observed before the other nodes of the system have had an opportunity to initiate reconfiguration, nodes in their quiescent states do not process messages with sequence identifiers different from the sequence identifier currently stored in local memory. In addition, some protocol messages are ignored during the quiescent state to prevent new protocols from executing. When the quiescent state ends and data structures are modified, no protocols are executing which makes the data structure changes much simpler and safer. For instance, proclaim, join, node connectivity, and group connectivity messages are all ignored during the quiescent state. Heartbeat messages are also not monitored so that nodes which have temporarily stopped transmitting heartbeats are not mistakenly considered not operational. Furthermore, messages of not yet completed protocols are transmitted with the previous configuration identifier and not with the reconfiguration sequence identifier. Heartbeat messages, on the other hand, are transmitted with the reconfiguration sequence identifier to help propagate the reconfiguration request. However, the heartbeat message is nevertheless accepted even when the sequence identifier does not match that of the receiver.

After the quiescent state has terminated 1035, as dictated by the period of time required for all of the executing protocols to terminate and for reconfiguration requests to be propagated to the other nodes of the system, each node implements the new configuration read from the configuration file, by performing data structure changes to reflect the change in topology, 930 and 1040. Subsequently, each node enters a grace period, 935 and 1050, as discussed below.

Because nodes of the system typically do not enter their quiescent states at the same time, the aforementioned grace period is used to allow all nodes to exit their quiescent states before any nodes resume normal operation. In this regard, the grace period, in one example, is equal to the length of the quiescent state. Furthermore, during the grace period, each node transmits messages with the new configuration identifier 940, begins honoring messages with the new configuration identifier (and initiating any new protocols as a result of these messages), and transmits protocol messages that normally have a limited amount of retries throughout the entire grace period 1055. Additionally, messages having configuration numbers different from the new

configuration number are ignored and heartbeat messages are not monitored, and no computation is made of the set of reachable nodes 1055.

Subsequently, after a period of time equal to the duration of the grace period has elapsed 945, normal operation is resumed (i.e., the set of reachable nodes is computed, any limits to the maximum number of retries is again enforced, and the monitoring of heartbeat messages is commenced, 950 and 1065).

In addition, FIG. 2 depicts a detailed example of a configuration file, which contains information regarding configuration of the system, and which is updated with a change in topology of the system. FIGS. 3-6B depict various examples of how a change in topology can be ascertained or implemented within a system.

Based upon the above discussion, as well as the further extensive discussion of change in topology and configuration files provided in the application, applicants respectfully request reconsideration and withdrawal of the 35 U.S.C. 112, first paragraph rejection.

Substantially, the Office Action contains a rejection of claims 1-2, 7-19, 24-37, & 42-52 under 35 U.S.C. 103(a) as being unpatentable over Moiin (U.S. Pat. No. 6,108,699) in view of Bertin et al.(U.S. Pat. No. 6,400,681). In addition, claims 3-6, 20-23 and 38-41 were rejected under 35 U.S.C. 103(a) as being unpatentable over Moiin and Bertin et al. in further view of Frank et al. (U.S. Pat. No. 6,532,494). These rejections are respectfully, but most strenuously, traversed and reconsideration thereof is requested.

Applicants' invention is directed to reconfiguring a network to reflect a change in the topology of the network by utilizing a predetermined length quiescent state. This quiescent state is entered by a node, upon receiving a reconfiguration request at the node. The node stays in the quiescent state for a period of time sufficient to allow at least one of the node to also enter a quiescent state. Upon termination of the quiescent state of the node, the node is reconfigured to reflect a change in topology of the network without checking with the at least one other node. Since Moiin and Bertin et al. (as well as Frank) do not teach or suggest the above-noted features, either alone or in combination, applicants respectfully request reconsideration of the obvious as rejection.

Initially, applicants note that both Moiin and Bertin et al. fail to disclose reconfiguring a node upon termination of a quiescent state, which is claimed by applicants in the independent claims presented. This aspect of applicant's independent claims does not appear to be addressed by the Office Action. At the bottom of page three, the Office Action does state that Moiin does not disclose reconfiguring a node to reflect a change in topology of the network without checking with at least one other node. However, this paraphrase of the applicants' claimed invention omits applicants' recited concept of proceeding with the reconfiguration after termination of the quiescent state at the one node. Since the Office Action does not address this feature of Applicants' independent claims, and since neither Moiin nor Bertin et al. discuss such a concept, applicants respectfully submit that this deficiency at least renders incomplete a rejection based on an alleged combination of Moiin-Bertin et al. For at least this reason, reconsideration of the obviousness rejections is requested.

Moiin describes a technique for modifying membership in a clustered distributed computer system and updating system configuration. In Moiin, each node receiving a reconfiguration message, which is referred to as a petitioned node, determines all other nodes which the node is connected to and responds with a reconfiguration message which proposes a respective new cluster including all such nodes. This petitioning and the petitioned nodes collect all reconfiguration messages, and if the reconfiguration messages unanimously propose the same proposed cluster, the proposed cluster is accepted as a new configuration. Unanimous agreement as to membership in the cluster is required by Moiin. In another aspect, multiple nodes can leave a cluster simultaneously. Failure to receive messages from a particular node in a predetermined period is detected as a failure of the node. In response to a detected failure, the node detecting the failure sends a reconfiguration message. Each node receiving the reconfiguration message broadcasts in response thereto a reconfiguration message to all other nodes and determines from which nodes a reconfiguration message is received back. Thus, each node determines which other node is operatively connected and configures a proposed new cluster which includes as members the connected nodes.

For an alleged teaching of various aspects of Applicants' claimed invention, the Office Action references column 2, lines 23-34 & 39-44 of Moiin. The applicability of this

discussion in Moiin, as well as the characterization thereof relative to applicants' invention is respectfully traversed. The Office Action cites these lines of column 2 of Moiin relative to applicants' recited language of claim 1, but provides no explanation as to how the language applies to applicants' recited functionality.

For example, applicants' independent claims recite: "...upon receiving a configuration request at one node of the plurality of nodes, entering a quiescent state at the one node, wherein the one node remains in a quiescent state for a predetermined period of time sufficient to allow at least one other node of the plurality of nodes to also enter a quiescent state ..."

Applicants recite functionality in their process which requires receipt of a reconfiguration request at a node of the plurality of nodes of the network. Upon receiving this reconfiguration request the node then enters a quiescent state. No similar functionality is discussed, suggested or implied by Moiin (nor the other art of record). In view of this deficiency, applicants' respectfully request reconsideration and withdrawal with the obvious rejection.

Applicants' further recite that the quiescent state entered by the node after receiving a reconfiguration request is for a sufficient period of time to allow at least one other node of the network to also enter a quiescent state. Again, no such functionality is described by any of the applied art. In Moiin, the predetermined period discussed at column 2 (as referenced in the Office Action) refers to a failure to receive messages from a particular node for a predetermined period. This failure to receive messages from a node is determined as a failure of the node. Clearly, the use of the predetermined period in Moiin is in a different process than that recited by applicants. Again, applicants' functionality recites first receiving a reconfiguration request at a node and second responsive to receipt of that reconfiguration request placing the node in a quiescent state. No similar functionality is taught, suggested or implied by Moiin, or the other applied art.

Further, Moiin fails to teach or suggest applicants' claimed element of upon determination of the quiescent state, reconfiguring the node to reflect the change in topology network without checking with at least one other node. (For example, claim 1.) Applicants

note that it is admitted at the bottom of page three of the Office Action that Moiin fails to disclose this element.

For the above reasons, applicants' respectfully submit that Moiin fails to teach or suggest multiple aspects of applicants claimed invention.

Bertin et al. fails to overcome the deficiencies of Moiin as applied against applicants' claimed invention. Bertin et al. is directed to a technique for minimizing the connection set up time in high speed packet switching networks. The Office Action references column 8, lines 34-44 of Bertin et al., which states:

The network topology information is updated when new links are activated, new nodes added to the network, when links or nodes are dropped, or when link loads change significantly. Such information is exchanged by means of control messages with all other Route Controllers to provide the up-to-date topological information needed for path selection (such database updates are carried on packets very similar to the data packets exchanged between end users of the network). The fact that the network topology is kept current in every node through continuous updates allows dynamic network reconfigurations without disrupting end users logical connections (sessions).

A careful review of the above-noted material from Bertin et al. fails to uncover any teaching, suggestion or implication of the above-noted deficiencies of Moiin when applied against the independent claims presented herein. Further, although there is discussion of maintaining network topology current in every node, there is no discussion, teaching or suggesting in Bertin et al. of: (1) receiving a reconfiguration request at one node; (2) responsive to receiving a new reconfiguration request entering a quiescent state at that node, wherein the quiescent state is a sufficient predetermined period of time to allow at least one other node of the network to enter a quiescent state; and (3) upon terminating the quiescent state at that one node then reconfiguring the one node to reflect the change in topology of the network without checking with at least one other node. Applicants respectfully submit that the Office Action does not address these concepts recited in the independent claims presented. For instance, neither patent teaches or suggests applicants' claimed element of upon receiving the reconfiguration request at one node of the plurality of nodes, entering a quiescent state at the one node. Bertin et al. does discuss maintaining topology of the

network current at a node, but does not explain how the updated network topology needs to be maintained.

Further, there is no discussion, teaching or suggestion in Bertin et al. (nor Moiin) of the node remaining in the quiescent state for a predetermined period of time sufficient to allow at least one other node of the plurality of nodes to also enter a quiescent state (e.g., claim 1). This aspect of the claimed invention is simply missing from Bertin et al. (and Moiin).

Yet further, there is no discussion, teaching or suggestion in Bertin et al. of reconfiguring a node to reflect the change in topology, upon termination of the quiescent state (e.g. claim 1). In Bertin et al., reconfiguration is performed to continuous updates.

In summary, applicants' respectfully submit that a careful reading of Bertin et al., and in particular, column 8, lines 34-44 (cited in the Office Action), fails to teach or suggest various features of applicants' claimed invention. Again, there is no teaching or suggestion of entering a quiescent state upon receiving a reconfiguration request at a node, of reconfiguring the node upon termination of the quiescent state, nor of reconfiguring the node without checking with the at least one other node as claimed by applicants.

Since both Moiin and Bertin et al. fail to teach or suggest multiple aspects of applicants' claimed invention, applicants respectfully submit that the combination of Moiin and Bertin et al. also fails to teach or suggest their claimed invention. Thus, applicants respectfully request an indication of allowance of the independent claims.

The dependent claims are believed allowable for the same reasons as the independent claims, as well as for their own additional characterizations. As noted, claims 3-6, 20-23 and 38-41 were rejected as obvious over Moiin in view of Bertin et al. and in further view of Frank et al. This combination also fails to teach or suggest one or more features of the applicants' claimed invention. Reconsideration and withdrawal of the objection is therefore requested.

As described in detail above, Moiin and Bertin et al. fail to describe, teach or suggest at least applicants' claimed features of entering a quiescent state upon receiving a

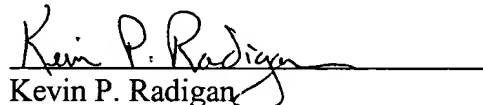
reconfiguration request at a node, and remaining in a quiescent state for a predetermined time sufficient to allow for at least one other node to also enter a quiescent state, and upon termination of the quiescent state at the node, reconfiguring the one node without checking with the at least one other node.

Frank et al. fails to overcome the deficiencies of Moiin and Bertin et al. when applied against applicants' claimed invention. For example, Frank et al. fails to describe or suggest use of a quiescent state when reconfiguring. Frank et al. describes at column 5, lines 18-20 that if a node fails to receive a heartbeat message from another node within a predetermined time interval, then a cluster enters reconfiguration mode. Clearly, this cited language in Frank et al. does not teach, suggest or imply applicants' above-summarized invention of the independent claims. The use of a predetermined time interval in Frank et al. is to ascertain whether a node becomes non-responsive, which is a different concept and process than that recited by applicants'.

Based on the foregoing, applicants respectfully submit that the dependent claims patentably distinguish over Moiin in combination with Bertin et al. and/or Frank et al. Reconsideration and withdrawal of the rejections based thereon is therefore requested.

In view of the above, allowance of all claims is respectfully requested. Should the Examiner wish to discuss this case with applicants' attorney, the Examiner is invited to contact applicants' attorney at the below-listed number.

Respectfully submitted,



Kevin P. Radigan
Attorney for Applicants
Registration No.: 31,789

Dated: September 02, 2003.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579